

General Disclaimer

One or more of the Following Statements may affect this Document

- This document has been reproduced from the best copy furnished by the organizational source. It is being released in the interest of making available as much information as possible.
- This document may contain data, which exceeds the sheet parameters. It was furnished in this condition by the organizational source and is the best copy available.
- This document may contain tone-on-tone or color graphs, charts and/or pictures, which have been reproduced in black and white.
- This document is paginated as submitted by the original source.
- Portions of this document are not fully legible due to the historical nature of some of the material. However, it is the best reproduction available from the original submission.

BOLT BERANEK AND NEWMAN INC
CONSULTING • DEVELOPMENT • RESEARCH

BBN Report No. 2993

A.I. Report No. 22

"ROBOT" COMPUTER PROBLEM SOLVING SYSTEM

Final Progress Report

Contract No. NASW-2572

Joseph D. Becker

E. William Merriam

23 October 1974

Prepared for:

National Aeronautics and Space Administration
Washington, D.C. 20546

(NASA-CR-142860) ROBOT COMPUTER PROBLEM
SOLVING SYSTEM Final Report (Bolt, Beranek,
and Newman, Inc.) 30 p HC \$3.75 CSCL 06D

G3/63

Unclass
17709

N75-25641

BBN Report No. 2993
A.I. Report No. 22

"ROBOT" COMPUTER PROBLEM SOLVING SYSTEM

Final Progress Report
Contract No. NASW-2572

Joseph D. Becker
E. William Merriam

23 October 1974

Prepared For:
National Aeronautics and Space Administration
Washington, D.C. 20546

INTRODUCTION

The following is a report on progress made on NASA contract number NASW-2572. Our work on this contract encompasses the conceptual, experimental, and practical aspects of the development of a Robot Computer Problem Solving System. We categorize the progress made in these three domains as follows:

- I. Conceptual Issues
Formulating the distinctive characteristics of the approach taken by our project in relation to various studies of cognition and robotics.
- II. Experimental Issues
Structuring of the vehicle and eye control systems and defining the information to be generated by the visual system.
- III. Practical Issues
Continuing support to the JPL robot project and disseminating information about our own project.

These topics are discussed in the following pages, using the same outline as that given above.

I. CONCEPTUAL ISSUES

A. Clarification of our Research Methodology

Our research on the development of a robot computer problem solving system is relevant to and incorporates ideas from several different fields of study: computer science, psychology, control theory, physiology, decision theory, and so forth. At the same time, our project is quite novel, and it differs greatly from the sort of research that is traditional to any one of these fields, e.g. "pure" control theory or standard cognitive psychology. Because our work is such an unusual combination of many disciplines, we have sometimes found it difficult to communicate our approach to professional people who are accustomed to thinking in terms of the traditional academic rubrics. The misunderstandings that have resulted on such occasions have been productive, because they have forced us to make explicit some of the assumptions about research methodology that we had formerly been working under without having a clear awareness of them. In this section we will indicate some of the ways that we have been able to characterize our approach, and to relate it to more conventional methods of research.

Animals vs. Machines

The object of our research is an elusive commodity that

we call "cognitive organization", i.e. those principles of mental functioning that allow a being to think. But what kind of thinking beings are we interested in, animals or machines? The answer is that in our approach, the two are inextricably linked. We feel that animals are necessary to the study of cognition because they provide the only working models of thinking that we know to exist; indeed, people generally use animal or human behavior to define cognition, and refuse to attribute the name "thinking" to unnatural problem solving behavior in machines (e.g. the rapid arithmetic of a desk calculator). On the other hand, we feel that animal-simulating machines, i.e. robots, are equally vital to the study of cognition, because they provide us with new paths of experimental investigation that are entirely unavailable to traditional psychology, as will be explained shortly. It goes without saying that thinking machines should also prove to be exceedingly useful, but practical applications are not the immediate goal of our particular project.

Observation vs. Introspection

Before there were machines with the potential to think, the study of cognitive organization was exclusively a branch of psychology. The two main avenues of investigation open to psychology are observation and introspection. Unfortunately, neither of these approaches can come very

near to the issues that concern us in this project, such as the representation of spatial experience and the choice of reasonable but not necessarily optimal actions. Direct observation of the external behavior of animals or humans cannot give reliable insight into cognitive organization, which after all is "in the head." But human introspection, aside from its many known pitfalls, has access only to "near-conscious" cognition; no adult can give a meaningful account of how he visually recognizes a chair as such, or how he knows what to do with his hand when his nose itches. Thus, the cognitive patterns that we are interested in are precisely those that are too thoroughly overlearned to be available to introspection.

Given that our object of study lies outside the reach of all traditional psychological methodology, our approach is simply to make the best use possible of whatever psychology has to offer us. We try to remain abreast of the literature of cognitive psychology, and we glean as much information as we can from our own introspections. Although these sources do not give us direct information about cognitive organization, they sometimes can be helpfully suggestive, and sometimes they indicate boundary conditions on the processes that we are looking for, even when they do not reveal the workings of the processes themselves. (For example, both observation and introspection make it obvious that the process of searching perceptual memory is highly

associative, and incredibly more efficient than any search process known to computer science; these facts are important, even though they tell us nothing about how animate memory search is actually performed.)

Analysis vs. Synthesis

As we have said, the traditional scientific methods of psychology are not adequate to reveal the secrets of cognitive organization, especially at a detailed level. If we wish to go beyond these analytic approaches, we really have no choice but to adopt a synthetic technique, i.e. to build our own cognitive system from scratch. The benefits of a synthetic approach are several. Its primary advantage is that the inner workings of any system that we build ourselves are of course known to us, so there is not the sort of hopeless inaccessibility that limits traditional psychology. There is the fact that "getting one's hands dirty" in building something gives one a much better intuition about its functioning than does mere passive observation or armchair contemplation. There is also the fact that a synthetic approach often forces a clear separation between the essential and the irrelevant aspects of a phenomenon. For example, mankind was never able to understand the flight of birds until he built his own flying machines and learned that the flapping of the wings was inessential to winged flight per se. He found that what was

essential were aerodynamic principles such as lift. With this new understanding, man was able to go back to the analysis of bird flight, and to work out the proper role of flapping, which of course is essential for birds. Similarly, we would hope that once we have built machines that can exhibit some cognitive behavior, we will have learned what are the basic principles of cognitive organization itself (analogous to aerodynamic principles), and what are the incidental aspects that constitute the "flapping" of the human mind. Finally, a by-product of the synthetic approach is that it can produce devices which have the potential for practical usefulness.

These advantages of the synthetic approach do not by any means imply that the information gained from analytic science should be ignored. Perhaps it should seem obvious that analysis and synthesis are to be joined in collaboration, not in competition, but unfortunately this obvious principle has often been violated in the history of machine intelligence. In the old days there was a sort of hubris that might be expressed as "If we can make machines that add better than humans, we can make machines that think better than humans!" This attitude has been considerably damaged in the collision with harsh realities. But a much more insidious phenomenon is the way in which many synthetic research projects degenerate into the production of a truly artificial system which has moved so far away from the

analytic facts of animate cognition that it has little scientific value, while at the same time its scientific pretensions make it too poorly engineered to be of practical value. We cannot afford to be too critical of this sort of "degeneration", however, because a fair amount of it is inevitable in the synthetic approach, as we shall now explain.

Purity vs Practicability

So far we have said that we attempt to combine the analytic study of cognitive organization in animals and man with the synthesis of our own mechanical cognitive system, our robot. We do strive as hard as we can to keep this combination "pure", in other words to attempt to make the robot really be a simulated animal and not a trumped-up artificial creation with no particular relationship to animate cognition. But there are two factors that carry us very far from absolute purity in this regard.

The first is simple practicability. For example, perceptual memory search in animals is implemented in brains containing millions or billions of neurons acting simultaneously. This undoubtedly accounts for the uncanny speed and efficiency of such search. Even if we understood the logic of this kind of search, it would be totally impossible to write a simulation program for present-day

computers that would run fast enough to be usable. Given the present state of computer science (which includes programming languages as well as hardware), we very often are simply forced to cut corners if we want our simulation to be runnable at all. The best we can do is to note the places where we have cut corners, and not attach any theoretical importance to them.

The second factor is more insidious. It comes from the fact that the computer programming of our simulation forces us to an extremely fine level of detail, whereas all the observational or introspective information available from psychology is at an extremely gross, general level. This means that when the time comes actually to program the simulation, our means of conceptualizing and implementing the detailed program are much more closely determined by the computer than they are by any information we may have about the process we are trying to program! One way of looking at this is to say that in closing the huge gap between the synthetic behavior of the program and the analytic information that we have as boundary conditions, we are forced to start from the synthetic side of the gap, which is most unfortunate. This problem is unavoidable, so all we can do is insist that our simulation be as psychologically realistic as we know how to make it. Also, we assume that we will have to rewrite the whole simulation over several times from scratch, with each new version being based on

ideas which became clear only through the eventual failure of the previous version to meet its psychological boundary conditions.

In summary, our approach to the investigation of cognitive organization includes paying attention to both animals and machines, making use of data from both observation and introspection, combining the advantages of both the analytic and synthetic methods, and keeping our system as pure as we can within the constraints of practicability. This eclecticism is even more difficult to carry out than it is to explain, but we feel that it is the only methodology which will lead us to an eventual understanding of the mechanisms of cognition.

B. Functioning of the Visual System

Considerable study has gone into the design of the visual system for our robot simulation, and the system is still under revision in many respects. Many of the problems that have influenced our design of the visual system could be considered to be "conceptual issues", but they are so closely related to practical considerations that we have decided to discuss them all together under "Experimental Issues" in Section II.C of this report.

II. EXPERIMENTAL ISSUES

A. Perspective on our Development of the World Simulation

As we suggested in Section I.A, the development of our robot simulation program has to proceed literally from the ground up: we first implement a Martian terrain environment, then the physical robot and its sensori-motor system, then the physics (forces and movement) in this world, then the primitive sensory operations ... and finally on to the perceptual and higher cognitive operations. In the past, we have tended to view this simulation as being made up of two distinct parts, whose boundary is indicated by the "..." in the previous sentence; that is, the environment, physics, motor, and sensory systems were considered to be the "World Simulation" and the perceptual and problem-solving systems were considered to be the "Cognitive Simulation".

As our work on the simulation has progressed, we have come to doubt if this kind of dichotomy can be maintained. There are two tendencies we have discovered which suggest that our simulation will be much more homogeneous than we had originally envisioned it. The first is that we are never able to find boundary lines among what we imagined to be the separate components of the system. The world interacts with the motor system, which interacts with the physics, which interacts with the sensory system, and so on.

"Sensation" is presumed to be physical, and "perception" is presumed to be cognitive, but certainly these two concepts are merged even in theory, so that in a practical simulation they become inseparable.

The second phenomenon that we have encountered is that when we discover a shortcoming at any level in the simulation (and the whole purpose of this research project is to discover the limitations and extensions of our current knowledge), to eliminate that shortcoming often requires revising much of the whole simulation. For example, we originally had limited the robot to a fixed, polygonal path through the environment. Later we decided that this, besides being unrealistic, did not give the robot sufficient freedom to make an interesting choice of actions, so we removed the idea of a restricted path. But then, we had to change the robot's motor system to allow it to move in curves as well as in straight segments. We also had to give physical properties to mountains, hills, and crater rims, so that the robot would bump into them and stop, rather than passing right through them as it would have in the initial simulation. But then, since it could bump into these edges, the robot had to be able to see the edges as well, which was not included in the original visual system. But then, since the mountains, hills, and so forth now had physical and visual reality, we had to solve the difficult problem of simulating the visual occlusion of one object by another

object in front of it.

Now, all of these extensive revisions to the "world simulation" arose because of a single change (dropping the fixed path), whose primary motivation was not physical realism, but cognitive realism (i.e. that the robot should have a lot of choices to make about where it might go). It is perhaps worth pointing out too that these changes arose purely from the theoretical considerations of our research, and in no wise from technical considerations of computer implementability. Indeed, these changes were discussed explicitly in last year's proposal to this contract ... where we asserted that we would not be making them because they were too difficult to program!

In summary, we have come to see our robot simulation as a single evolving entity whose various aspects are closely interrelated. Since we expect this to be an experimental program where we will often try out ideas that turn out to be unsuccessful, we must expect in consequence that we will be continually revamping our simulation, even those "lower-level" aspects such as the "world simulation" which had seemed to reach a stable level of development. Needless to say, the kinds of changes we will be making will seldom involve scrapping whole chunks of our simulation and reworking them de novo; rather, they will be revisions and additions in the direction of greater realism or

sophistication. The descriptions of our current work in the remainder of this section clearly show that we have put in much that is new, while retaining many of our original ideas.

B. Improvements to the Simulation

Most of our previous work on robot problem solving has centered on activities which relate to the visual perception of objects. Our work on tracking an object has been the single exception to this, since it was necessary to coordinate the motion of the robot's body with the motion of its eye. Because our interest at that time was primarily in the perceptual aspects of robot problem solving, we chose to restrict the motion of the robot to a well-defined path so that we could temporarily ignore both the conceptual and practical problems of implementing a complete motor control system. As a result of this simplification, we have come to a point in our understanding of robot problem solving where we are ready to allow the robot to set out on its own, making decisions about where it is to go. In order to allow it to do this, we completely redefined the robot's motion control system which, paradoxically, got us embroiled right back in issues of visual perception, as described in Section II.A. Thus, in this section we describe the newly completed robot body and eye motion control systems and in the next section we discuss areas that are only partially

complete which involve the type of information which is to be generated by the visual system.

Vehicle Motion Control

Before we could specify the physics of a vehicle control system, it was necessary to define what kinds of actions we wanted the vehicle to perform. We decided upon a set of essentially car-like properties that allow the robot to go forward and backward at any chosen speed (but within specified limits) and to turn in both directions. The single non-car-like property is that we wanted the robot to be able to turn in place about its center point. The reason for this is not so much conceptual in nature, but rather is a result of the mathematics for simulating the path taken by the robot. (Here is a good example of the Purity vs. Practicability dilemma mentioned in Section I.A. In order to make the simulation run in a reasonable amount of time, we chose to approximate the robot's (conceptually) curved path with a sequence of short straight line segments. But doing this could result in some anomalies when operating very close to the edge of a terrain object. Since these anomalies could get us into trouble with the cognitive system later on, we chose to eliminate them by making the robot's motion a turn-then-move operation, which leads to the turn-in-place characteristic of the simulated vehicle.)

The parameters to the vehicle control system consist of two velocities -- a forward velocity and a turning velocity. Either velocity can have positive or negative values to indicate direction. The interpretation of the velocities is that they are desired velocities. The vehicle control system will use these velocities to alter the behavior of the vehicle within the limits allowed. For instance, in order to allow for the effects of a bound on maximum acceleration and deceleration without introducing the clumsiness of control using acceleration, the rate of change of velocity is limited. So also is the maximum positive and negative velocity that is allowed to be reached. Thus, if the vehicle is going forward (a positive velocity) and the control system is given a negative velocity as a parameter, the vehicle will gradually slow down, reverse direction, then gradually speed up until it reaches the desired negative velocity; unless the parameter exceeds the maximum negative velocity allowed, in which case the vehicle will maintain that maximum and no more.

Another aspect of the vehicle motion control system which we have programmed is that of feedback of information to higher-level systems. For instance, if the vehicle has bumped into an object, that fact should be reported along with information about where on the robot the hit took place. So far, we are uncertain (both from a conceptual and practical viewpoint) as to what vehicle control feedback

information will be important later on. Therefore, we are making all information available that we think might be useful. This includes: the actual X and Y position of the robot in the world; the vehicle's actual forward and turning velocity; and whether the robot has hit or crashed into any object.

Notice that we have started to develop a hierarchy of control. The vehicle takes "turn-then-move" commands from the control system. The control system takes desired velocities from some higher level system not yet completely defined, but which we view as the sensori-motor system. We envision the sensori-motor system as receiving commands such as "go faster" from the cognitive system. Thus, our work on vehicle control is more than simply a redefinition of how the vehicle will work, but it is a synthesis of the conceptual ideas generated during the earlier phases of our project. A similar phenomenon holds for the visual system.

Eye Motion Control

As a result of our conceptual work, the visual system has changed in detail somewhat (See Section II.C), but the control parameters themselves have remained the same. These are: the angular position of the eye relative to the center line of the robot; the focal distance; the angular width of vision; and the depth of vision. The last two parameters

are specified by simply giving a value, and changes are assumed to be instantaneous.

The first two parameters are more complex. They can be specified by providing either positional information (as with the last two parameters) or by specifying a (positive or negative) velocity which the eye is to achieve. The reason for the two modes of control is that there are cases where each seems conceptually correct. Sometimes it is desirable to look in a particular direction and for this, positional control is required. Other times, such as in the tracking of an object, it is desirable to be able to move the eye at a certain speed and for this, velocity control is used. Regardless of which mode of control is used, however, the repositioning and refocusing of the eye is limited by a maximum velocity and maximum change in velocity in a manner similar to the parameters controlling the vehicle. Thus, if the eye is stopped in a position to the far left of the robot and the positional parameter is specified for the far right, this change will not take place instantaneously, but rather at a pace determined by the established limits.

As with the vehicle system, certain feedback information will be returned which indicates the current state of the eye. Much more important though is the visual information actually returned by the eye. We at one time had this problem solved, but as a result of redefinition of

the visual system, the exact information returned by the eye is still an open question which we will now discuss.

C. Work in Progress: The Visual System

The work we are doing at present centers around redesigning the robot's visual system. As we pointed out in Section II.A, an aspect of our simulation often reveals avenues for improvement after only a limited amount of testing, and in this spirit we are incorporating many new ideas into the revised visual system based on our experience with the old one.

Geometry of the Visual Field

The visual field has been redefined to include a central region and a peripheral region. The peripheral region is essentially the region within which lateral, near, and far vision takes place, and the central region is where foveal vision takes place. The peripheral region spans an area approximately 70 degrees to each side of the direction in which the robot is looking, and from a few inches in front of the robot to the horizon. The central region has essentially the same external geometry as our old complete eye and is located along the center line of the peripheral region and completely within it. The eye control parameters discussed in Section II.B all apply to the central region

and do not effect the peripheral region except to change the position of its center line. Aside from rotating along with the eye, its size does not change.

The reasons for the change to the geometry of the visual field stem from the fact that certain anomalies arose with the old system which were a direct result of the peripheral fields becoming smaller as the eye focused down on an object. It was possible, for instance, to be focused on a distant mountain and not see a nearby hill that was in the direct path from the robot to the mountain. It was also possible to be so focused down on an object that the peripheral field of view was essentially non-existent. No such anomalous cases exist with the new visual system.

We will now discuss the sensory aspect of the eye, i.e. what it sees, and we will concern ourselves only with vision in the central region. The visual properties of the peripheral field remain to be worked out, but are sure to be quite simple compared with the finely-structured sensory information brought in by the central field.

Levels of Processing

If we examine the physiology of vertebrate vision, we find that it is quite fallacious to imagine that a single sensory entity, "the eye" sends some simple commodity called "visual information" to a single entity, "the brain".

Rather, we find that there are many distinct processing stations which operate on the visual image, starting with more than one integrative layer in the retina itself. Thus, there is no single answer to the questions "What does the eye see?" or "What visual information does the brain receive?", but rather there are many successive levels of assimilation, starting at the retina and ending in the cortex.

If we look at vision psychologically rather than physiologically, our conclusion is just the same. Psychologists sometimes for convenience consider "sensations" to be light-meter-like readings sent by the eye, and "perceptions" to be analytical interpretations of the sensations as made by the brain, but in truth this distinction is fictional. All we really know is that light rays go into the eyes, and cognitive interpretations arise in the brain, and that a lot of processing has gone on in between.

This conclusion applies to vision in our robot simulation model, and it raises difficulties for us. We cannot simply have a program called "the eye" which takes simulated light as input and gives simulated "vision" as output -- because this would be unrealistic, and probably unworkable as well. This means that we must set up an explicit hierarchy of visual processing. This is an onerous

task because we have very little to use as design criteria beyond intuition: psychology is of little help because the psychology of visual processing is not understood in the kind of detail that our simulation demands, and physiology is of little help because our robot's eye is after all artificial, and its primary operations are not at all like those of the vertebrate retina.

As matters stand, we have not yet decided upon the precise levels of visual processing to be used in our model. This fact will lend a certain vagueness to the rest of the discussions in this section, since we are as yet undecided as to what processing will be done where. All such decisions will of course be forced by the process of programming the new visual system.

The Seeing of Point Objects

In previous reports we have noted that the robot's Martian environment contains a large number of rocks and "unidentified objects", which we collectively call point objects. We have explained that we have put point objects in the robot's world in order to experiment with the recognition of a scene on the basis of constellations of individual objects, where the process of perceiving each individual object is not particularly interesting, i.e. is particularly simple. What is this simple process of seeing

point objects?

Briefly, a point object is seen as a list of attribute-value pairs associated with a geometric point location. Intuitively speaking, a rock might be seen as "Color-Gray at R-Theta, Texture-Rough at R-Theta, ...". In the old visual system, the polar coordinates of the point object with respect to the eye were not reported directly, but rather were sorted into one of 15 visual subfields. We have felt this pre-sorting to be an undue a-priori restriction on the input, and in the new visual system, the R-Theta coordinates themselves will be reported, at least at the lowest levels. We are considering introducing pseudo-random error into the radial coordinate, since vertebrate vision is not very precise as to ranging. (In angular measurement, at least relative measurement, the human eye is so accurate that we may consider it errorless.)

As in the old system, not all the possible attribute-value pairs of a point object will necessarily be seen at once. Several factors will determine what is seen at any given time. Different attributes have different visibilities (e.g. it is easier to see a rock's color than its texture). Different values sometimes have different visibilities (e.g. it is easier to see a red rock than a gray one). Near objects are easier to see than far ones. Objects near the very center of the visual field (the fovea)

are much easier to see than those near the periphery. Objects or attribute-pairs that the system is intentionally looking for are easier to see than those which are received by accident (this feature will be a new addition to the visual system). Also, there is a notion of "focusing" or "concentration" which allows one object to be seen in more and more detail, while all others are seen less and less fully. This factor, which in the current system is part of the initial selection of which attribute-value pairs are seen, may in the new system be moved to a later level of processing, since it is more a mental operation than a visual one. Finally, we are considering including a probability-directed pseudo-random choice as to which attribute-value pairs are seen at any given time, since the same object never really looks the same in any two encounters, no matter how similar the visual circumstances may be.

We might mention that the actual attribute-value pairs we will use will not be semantically meaningful, such as "Color-Gray", but rather will be simple codes with numerical values, such as "B-6". This sort of coding is realistic, inasmuch as the signals involved in the lower levels of our visual systems certainly are quantitative, and certainly do not involve English words. In our model, we can experiment with various spatial and probability distributions in the assignment of values to the attributes of individual point

objects. This can give rise to various situations of psychological interest, e.g. a constellation of rocks all of more-or-less the same "F-attribute", or a single unidentified object whose "D-attribute" value is markedly different from all the rest.

So much for the considerations that enter into the "simple" case of point object vision!

The Seeing of Terrain Objects

Along with point objects, the robot's environment contains a landing module (formerly a point object), craters, hills, and mountains, all of which are collectively known as terrain objects. Terrain objects are polygonal, and their vertices are essentially a class of point objects, so we will not concern ourselves primarily with the seeing of the vertices in this discussion. What concerns us here is the manner in which terrain objects are to be seen as connected wholes, i.e. (in effect) the seeing of the edges which join the vertices into connected, spatially-extended structures. We will ignore here the question of whether the edges of terrain objects have attribute-value pairs of their own (a matter which should be easily handleable by analogy with point objects), and concentrate on the simple question of how the eye might see whether or not two terrain vertices are in fact joined by an edge, i.e. whether or not they are

successive vertices of the same object. This question, though simple to state, gives rise to many complexities, and we are only part way toward answering it fully.

Our first assumption is that there is some range R such that if the two points are within distance R of the robot, the eye can always sense their true connectivity (i.e. whether they are joined by an edge or not), just as though connectivity were a directly-visible property. This assumption models the fact that vertebrate visual systems can make use of a wide range of subtle spatial clues, such as textural gradients, shadows, parallax, and so forth, so as to correctly see the connectivity of two points. Rather than trying to introduce all of these fine details into our simulation, we will simply sum up their net result, which is that connectivity is effectively a directly-visible property at close range. We are not saying that this range R is a constant, but simply that for any two point objects one can compute an R , and if they fall within it, perception of connectivity ceases to be a problem.

Outside of R , it remains a problem. Doubtless the connectivity of distant points cannot be detected with certainty, but rather must be reported in terms of confidence values, e.g. 0-100%. This means that we now must deal not only with "sensations" but also with "sensory hypotheses" (this problem arises, but much more weakly, in

the case of error in the detection of the radial position of point objects). At what level of visual processing are such hypotheses to be stored, checked out, or revised? This we have not yet worked out.

The criteria on which the visual system might hypothesize that two terrain vertices are connected are at least:

1. Spatial proximity
2. Similarity of features
3. Sighting of an edge that might join them

(We do not yet even have a potential list of the criteria on which the system might hypothesize the two vertices to be dis-connected; this is a trickier problem, which might even involve higher cognitive processing.) The first two of these criteria can be judged by closeness measures which should be relatively straightforward. We will only say further about them that if we introduce pseudo-random error into the initial detection of the vertices, then these measures may lead to erroneous conclusions. This is a valid simulation of the fact that, for example, two overlapping mountains on the horizon might appear to run together into one, since the distance separating them might not be detectable at that range.

The problems come in criterion 3, "sighting of a edge that might join" the two vertices in question. First of all, we have not yet fully decided on the visual properties

of edges, but certainly their spatial orientation is not directly sensible. This fact can lead to situations in which the robot cannot visually distinguish between two similar terrain configurations, e.g. cannot see whether a gap is really a pass, or just a cove. This sort of ambiguity evidently is perfectly appropriate if the robot is sufficiently far away from the terrain in question. Our problems are how the robot should report a sensory ambiguity, how it should store it, and at what level of processing it should work at resolving the ambiguity.

In sum, the major difference between our new visual system and the old is that the new version will address the problem of perceiving terrain objects as connected surfaces. This is certainly an interesting issue in perceptual psychology, and an important facet of the robot's ability to comprehend its environment.

III. PRACTICAL ISSUES

Our work with the JPL Robot Project has continued with our assuming a consulting role on the issues of path planning algorithm, ground system development, and overall system design. During several trips to JPL we have been pleased to discover that many of our recommendations have been implemented and that the project is showing definite accomplishment. Our most recent trip to JPL coincided with the International Telemetering Conference, where we presented the paper "Experientially Guided Robots". This paper is a summarization of the work we have done on this contract including a list of some of the interesting unsolved problems. It was chosen as one of ten papers to be published in the Telemetry Journal as representative of the conference.